Assignee: Intel Corporation
Docket No.: 2207/7240

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANTS	:	Arlin R. Davis
SERIAL NO.	:	09/397,850
FILED	:	September 17, 1999
FOR	:	METHOD OF READING A REMOTE MEMORY
GROUP ART UNIT	:	2157
EXAMINER	:	Barbara N. Burgess

M/S: APPEAL BRIEF - PATENT
COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, VA 22313-1450

ATTENTION: Board of Patent Appeals and Interferences

APPEAL BRIEF UNDER 37 CFR 41.37

Dear Sir:

This brief is in furtherance of the Notice of Appeal, filed in this case on September 11, 2006.

1. REAL PARTY IN INTEREST

Intel Corporation is the real party in interest for all issues related to this application.

2. RELATED APPEALS AND INTERFERENCES

There are no other appeals, interferences, or judicial proceedings known to Appellant or Appellant's legal representative, which may be related to, directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

3. STATUS OF THE CLAIMS

This application currently contains claims 1-24. Claims 1-2, 9-10 and 17-18 were rejected under 35 U.S.C. § 102(e) as being anticipated by Talluri et al., U.S. Patent No. 5,884,313. Claims 3, 4, 11, 12, 19 and 20 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Talluri in view of Osborne, U.S. Patent No. 6,078,733. Claims 5, 6, 13, 14, 21 and 22 are rejected under 35 U.S.C. § 103 (a) as being unpatentable over Talluri in view of Osborne and further in view of Krishnan et al., U.S. Patent No. 4,922,416. Claims 7, 8, 15, 16, 23 and 24 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Talluri in view of Osborne and further in view of Krishnan et al. and Chow et al, U.S. Patent No. 6,052,387.

4. STATUS OF AMENDMENTS

There are no currently outstanding amendments. The attached listing of claims (section 8), reflects the status of the claims including this amendment.

5. SUMMARY OF THE INVENTION

According to the embodiments of the present invention, a message sent from a local device to a remote device may identify the message as a type of remote Direct Memory Access read operation. If the remote device determines that that is the case, the remote device is then to perform a remote Direct Memory Access (rDMA) write operation to the local device. In some situations, embodiments of the present invention provide an improved transfer of data between devices (e.g., in a network server system).

In the embodiment of claim 1, a method is presented including a first operation of sending a message from a local device to a remote device, via a network, the message including a transport header indicating a message type (see, e.g., page 14, lines 7-12 and Fig. 5, "Transport Header"; page 25, lines 8-15 and Fig.10). In a next operation at the remote device, it is determined whether the transport header of the message identifies the message as a remote Direct Memory Access (rDMA) read operation (see, e.g., page 25, lines 15-16). In another operation, a remote Direct Memory Access (rDMA) write operation is performed at the local device in accordance with data elements included in the message, if the transport header of the message identifies the message as said remote Direct Memory Access (rDMA) read operation (see, e.g., page 14, lines 12-15 and Fig. 6; page 25, lines 15-19 and Fig. 10).

In the embodiment of claim 9 a network device is provided for initiating a method to read data in a remote memory of a remote device directly into a local memory. The network device includes a network interface controller (NIC) (see, e.g., page 19, lines 3-6 and element 18 of Fig. 9). The network interface controller is configure to receive a message from the remote device, via a network, said message including a transport header indicating a message type (see, e.g., page 14, lines 7-12 and Fig. 5, "Transport Header"; page 25, lines 8-15 and Fig.10). The

network interface controller is to process the message to determine whether the transport header of said message identifies the message as a remote Direct Memory Access (rDMA) read operation (see, e.g., page 25, lines 15-16) and perform a remote Direct Memory Access (rDMA) write operation in accordance with data elements included in said message, if the transport header of said message identifies the message as said remote Direct Memory Access (rDMA) read operation (see, e.g., page 14, lines 12-15 and Fig. 6; page 25, lines 15-19 and Fig. 10).

The embodiment of claim 17 is directed to a tangible medium storing a plurality of program instructions, which, when executed by a processor installed in a network device causes the network device to perform the method described in claim 1.

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. The rejection of claims 1-2, 9-10 and 17-18 under 35 U.S.C. § 102(e) as being anticipated by Talluri et al., U.S. Patent No. 5,884,313 ("Talluri").

B. The rejection of claims 3, 4, 11, 12, 19 and 20 under 35 U.S.C. § 103(a) as being unpatentable over Talluri in view of Osborne, U.S. Patent No. 6,078,733 ("Osborne").

C. The rejection of claims 5, 6, 13, 14, 21 and 22 under 35 U.S.C. § 103 (a) as being unpatentable over Talluri in view of Osborne and further in view of Krishnan et al., U.S. Patent No. 4,922,416 ("Krishnan").

D. The rejection of claims 7, 8, 15, 16, 23 and 24 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Talluri in view of Osborne and further in view of Krishnan et al. and Chow et al, U.S. Patent No. 6,052,387. ("Chow")

7. ARGUMENT

Under 35 U.S.C. §102(b), a claim is invalid if the invention claimed therein is described in a patent issuing more than one year prior to the filing of the subject patent application. Though a patent reference may have issued early enough (or filed early enough as the case for 35 U.S.C. §102(e)), that reference must also enable one skilled in the art to practice the claimed invention. *See* Akzo N.V. v. U.S. Int'l Trade Comm'n, 1 U.S.P.Q.2d (BNA) 1241, 1245 (Fed. Cir. 1986).

Absent anticipation it may be possible to combine two or more patents together to render a claimed invention obvious, and unpatentable, under 35 U.S.C. § 103(a). In determining whether the claims are unpatentable it is necessary to look to what the references actually teach. "It is impermissible within the framework of § 103 to pick and choose from any one reference only so much of it as will support a given position, to the exclusion of other parts necessary to the full appreciation of what such reference fairly suggests to one of ordinary skill in the art." In Re Wesslau, 147 U.S.P.Q. (BNA) 391, 393 (C.C.P.A. 1965). Accordingly, a prior art reference must be considered in its entirety, and portions thereof must be taken in proper context. MPEP § 2141.02; Bausch & Lomb, Inc. v. Barnes-Hind, Inc., 230 U.S.P.Q. (BNA) 416, 419 (Fed. Cir. 1986).

Claims 1, 9 and 17

Independent claim 1, for example, recites: 1. Determining, at a remote device, whether the transport header of a message identifies the message as a remote Direct Memory Access (rDMA) read operation, and 2. performing a rDMA write operation at a local device in accordance with data elements included in the message, if the transport header of said message

identifies the message as the rDMA read operation. Similar features are found in independent claims 9 and 17. These features are not described or suggested by the Talluri reference.

Talluri refers to a system and method for efficient remote disk I/O access. As described at Col. 3, lines 49-53, and Fig. 5, Node B is to read information from a disk drive at Node A. As indicated at Col. 6, line 60 to Col. 7, line 2, the request message sent by Node B to Node A specifies the following information: the destination node (i.e., Node A), the source node (i.e., Node B), the information source (i.e., the location in the disk drive of Node A being read), and the global address to which the requested information is to be transmitted. There is nothing in this information that identifies the request message as an rDMA read operation, an express limitation in the claim. Moreover, no such identification is suggested by Talluri especially since "rDMA" (a term of art) is not even mentioned in the disclosure of Talluri.

In the Final Office Action, the "Response to Arguments" section provides the following argument,

"Talluri explicitly discloses the server receiving a remote data read request and the server performing a DMA operation. Talluri further references that a remote write operation can be performed in the same manner except for a minor change of the request message now indicates the write operation to be performed instead of the read operation. Therefore, Talluri indeed identifies the request as a rDMA read operation (column 3, lines 45-65, column 4, lines 9-11, 35-40, column 5, lines 1-13, 15-20, column 8, lines 5-8, 35-45, column 9, lines 26-31)."

The cited sections of Talluri are addressed below.

Col. 3, lines 45-65 is taken from the Background section of Talluri. In this section it is described how Node B is to read data from Node A. An IMMU entry is made in a buffer at

Node B that maps a physical address range for Node B with a global address range. Accordingly, when sending the read request message to Node A, the global address (relevant to Node B) is included in the message (see Col. 4, lines 5-6).

Col. 4, lines 9-11 and lines 35-40 are also taken from the Background section of Talluri. The read request message is received by Node A, and the server (i.e., Node A) sets up an OMMU entry which maps a global address range with a physical address range for the Node A server. To transfer the data from Node A to Node B, the data needs to be provided to NIC 60 (Fig. 1). This may be done using two DMA operations. The first (lines 21-25) is a DMA transfer from the disk/disk controller (elements 66 and 64 in Fig. 1) to local memory (RAM 56 in Fig. 1). The second (lines 32-37) is a transfer from local memory to the communications interface (NIC 60 in Fig. 1).

Col. 5, lines 1-13 and lines 15-20 refer to the Summary of the Invention section of Talluri. The operation for transferring data between Nodes A and B is much the same as described in the Background section discussed above. A key difference is that only one DMA transfer takes place, namely between the disk controller 55 and the NIC 60 (see lines 23-26).

Col. 8, lines 5-8 and 35-45 refer to the Description section of Talluri and further elaborate the Summary section. Accordingly, the DMA operation referred to at Col. 8, line 3 is referring to the DMA transfer from the disk controller of Node A to the network interface of Node A (see Col. 7, lines 40-43). Once the data is received by the network interface of Node A, it can then be transmitted to Node B.

Col. 9, lines 26-31 explicitly refer to a write operation as mimicking the read operation described in Fig. 7 and discussed above.

The references to DMA in Talluri are exclusively to transfers within components of a single Node (e.g., a server). With respect to independent claims 1, 9 and 17, there is no description or teaching in Talluri that a message received at Node A is processed to determine whether the transport header identifies the message as a type of rDMA operation. There is no description or teaching in Talluri that Node A performs an rDMA operation with Node B.

Since features of claims 1, 9, and 17 are neither shown nor suggested by Talluri, the rejection of these claims and 2, 10, and 18, which depend therefrom under 35 U.S.C. § 102(e) is in error. Appellant respectfully requests reversal of this rejection.

Claims 3-8, 11-16, and 19-24

These dependent claims provide limitations that further define the independent claims. The rejection of these claims relies on the disclosure of Talluri, which, as described above, fails to teach several features of the independent claims. The Osborne, Krishnan and Chow references fail to make up for the deficiencies of Talluri. Osborne refers to message processing. Contrary to the assertions made in the Office Action, there is no disclosure in the cited section of Osborne concerning the handling of rDMA read or write messages in a Virtual Interface or otherwise. It appears from the text of the Office Action that the Krishnan reference has been cited to show a last data segment and completion of a rDMA request and that Chow shows a last buffer bit. These three references taken individually or in combination with Talluri fail to teach or suggest determining whether a message indicates a rDMA read operation and performing a rDMA write operation in accordance with data elements included in the message as called for in

each of the pending claims. Accordingly, the rejection of claim 3-8, 11-6, and 19-24 under 35 U.S.C. §103(a) is in error. Appellant respectfully requests reversal of this rejection.

CONCLUSION

Appellants respectfully request that the Board of Patent Appeals and Interferences reverse the Examiner's decision rejecting claims 1, 3-11, and 13-21 under 35 U.S.C. § 103(a) direct the Examiner to pass the case to issue.

The Commissioner is hereby authorized to charge the appeal brief fee of $500.00 and any additional fees which may be necessary for consideration of this paper to Kenyon & Kenyon Deposit Account No. 11-0600. A copy of this sheet is enclosed for that purpose.

Respectfully submitted,

Date: January 16, 2007

Shawn W. O'Dowd
(Reg. # 34,687)

KENYON & KENYON LLP
1500 K Street, NW
Suite 700
Washington, DC 20005
(202) 220-4200 telephone
(202) 220-42501 facsimile
DC1-644012

APPENDIX

(Brief of Appellants Arlin R. Davis
U.S. Patent Application Serial No. 09/397,850)

8. CLAIMS ON APPEAL

The claims in their current form (including those claims under appeal) are presented below:

1. A method comprising:

sending a message from a local device to a remote device, via a network, said message including a transport header indicating a message type;

determining, at the remote device, whether the transport header of said message identifies the message as a remote Direct Memory Access (rDMA) read operation; and

performing a remote Direct Memory Access (rDMA) write operation at the local device in accordance with data elements included in said message, if the transport header of said message identifies the message as said remote Direct Memory Access (rDMA) read operation.

2. The method as claimed in claim 1, wherein the data elements in said rDMA read message identify a set of source buffers in the remote device which reference a remote memory in the remote device and a set of destination buffers in the local device that reference a local memory in the local device.

3. The method as claimed in claim 2, wherein the source buffers and destination buffers are registered with a Virtual Interface (VI) network interface controller of the remote device and the local device, respectively.

4. The method as claimed in claim 3, wherein the data elements of the rDMA read message specify the source buffers and destination buffers as multiple data segments with offsets and designate a channel of the Virtual Interface (VI) as a data path for the rDMA write operation.

5. The method as claimed in claim 4, wherein one data element of the rDMA read message specifies a last data segment and completion of the rDMA read request.

6. The method as claimed in claim 5, wherein the data is read from the remote memory of the remote device directly into the local memory of the local device over the Virtual Interface (VI), without making an intermediate copy of the data.

7. The method as claimed in claim 6, wherein the remote device builds virtual interface rDMA write descriptors with a sequence inserted into an immediate data field on a last data segment of each rDMA read request.

8. The method as claimed in claim 7, wherein the completion of the data transfer is processed at the local device, based on the immediate data that arrives with the last data segment of each rDMA write operation by the remote device.

9. A network device initiating a method to read data in a remote memory of a remote device directly into a local memory, said network device having a network interface controller (NIC) configured to perform the following:

receiving a message from the remote device, via a network, said message including a transport header indicating a message type;

processing said message to determine whether the transport header of said message identifies the message as a remote Direct Memory Access (rDMA) read operation; and

performing a remote Direct Memory Access (rDMA) write operation in accordance with data elements included in said message, if the transport header of said message identifies the message as said remote Direct Memory Access (rDMA) read operation.

10. The network device as claimed in claim 9, wherein the data elements in said rDMA read message identify a set of source buffers in the remote device which reference a remote memory in the remote device and a set of destination buffers in the local device that reference a local memory in the local device.

11. The network device as claimed in claim 10, wherein the source buffers and destination buffers are registered with the network interface controller (NIC) of the remote device and the network device, respectively.

12. The network device as claimed in claim 11, wherein the data elements of the rDMA read message specify the source buffers and destination buffers as multiple data segments with offsets and designate a channel of a Virtual Interface (VI) as a data path for the rDMA write operation.

13. The network device as claimed in claim 12, wherein one data element of the rDMA read message specifies a last data segment and completion of the rDMA read request.

14. The network device as claimed in claim 13, wherein the data is read from the remote memory of the remote device directly into the local memory of the network device over the Virtual Interface (VI), without making an intermediate copy of the data.

15. The network device as claimed in claim 14, wherein the remote device builds rDMA write descriptors with a sequence inserted into an immediate data field on the last data segment of each rDMA read request.

16. The network device as claimed in claim 15, wherein the completion of the data transfer is processed based on immediate data that arrives with the last data segment of each rDMA write operation by the remote device.

17. A tangible medium storing a plurality of program instructions, which, when executed by a processor installed in a network device causes the network device to perform the following:

receiving a message from a remote device, via a network, said message including a transport header indicating a message type;

processing said message to determine whether the transport header of said message identifies the message as a remote Direct Memory Access (rDMA) read operation; and

performing a remote Direct Memory Access (rDMA) write operation in accordance with data elements included in said message, if the transport header of said message identifies that the message is said remote Direct Memory Access (rDMA) read operation.

18. The tangible medium as claimed in claim 17, wherein the data elements of the rDMA read message identify a set of source buffers in the remote device which reference a remote memory and a set of destination buffers in the network device that reference a local memory.

19. The tangible medium as claimed in claim 18, wherein the source buffers and destination buffers are registered with network interface controller (NIC) of the remote device and the network device, respectively.

20. The tangible medium as claimed in claim 19, wherein the data elements of the rDMA read message specify the source buffers and destination buffers as multiple data segments with offsets and designate a channel of a Virtual Interface (VI) as a data path for the rDMA write operation.

21. The tangible medium as claimed in claim 20, wherein one data element of the rDMA read message specifies a last data segment and completion of the rDMA read request.

22. The tangible medium as claimed in claim 21, wherein the data is read from the remote memory of the remote device directly into the local memory of the network device over a Virtual Interface (VI), without making an intermediate copy of the data.

23. The tangible medium as claimed in claim 22, wherein the remote device builds virtual interface rDMA write descriptors with a sequence inserted into an immediate data field on the last data segment of each rDMA read request.

24. The tangible medium as claimed in claim 17, wherein the completion of the data transfer is processed based on the data that arrives with the last data segment of each rDMA write operation by the remote device.

9. EVIDENCE APPENDIX

No further evidence has been submitted with this Appeal Brief.

10. RELATED PROCEEDINGS APPENDIX

Per Section 2 above, there are no related proceedings to the present Appeal.